

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

This Page Blank (uspto)



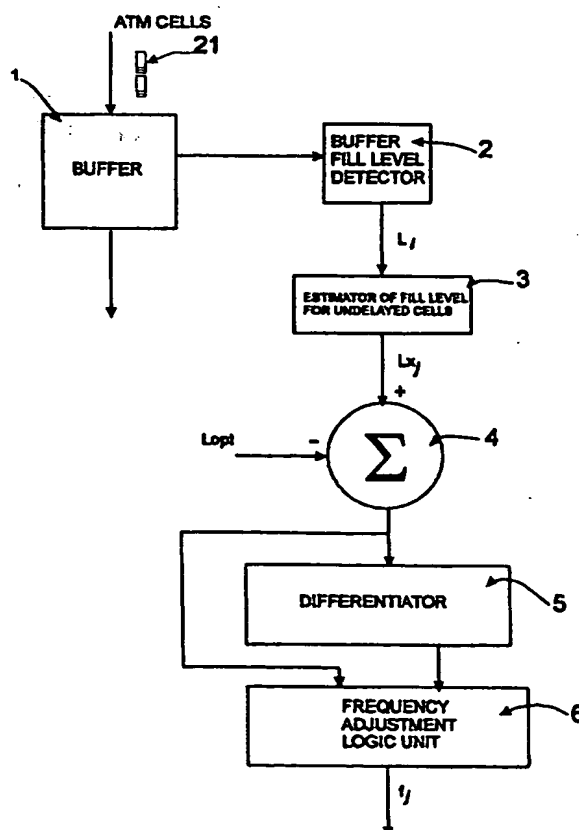
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04J 3/06	A1	(11) International Publication Number: WO 95/33320 (43) International Publication Date: 7 December 1995 (07.12.95)
<p>(21) International Application Number: PCT/CA95/00320</p> <p>(22) International Filing Date: 30 May 1995 (30.05.95)</p> <p>(30) Priority Data: 9410943.6 1 June 1994 (01.06.94) GB</p> <p>(71) Applicant (for all designated States except US): NEWBRIDGE NETWORKS CORPORATION [CA/CA]; 600 March Road, Kanata, Ontario K2K 2E6 (CA).</p> <p>(72) Inventor; and (75) Inventor/Applicant (for US only): COX, Neil [CA/CA]; 6158 Elgin Avenue, Burnaby, British Columbia V5H 3S4 (CA).</p> <p>(74) Agent: MITCHELL, Richard, J.; Marks & Clerk, P.O. Box 957, Station B, Ottawa, Ontario K1P 5S7 (CA).</p>	<p>(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SI, SK, TJ, TT, UA, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, MW, SD, SZ, UG).</p> <p>Published With international search report.</p>	

(54) Title: CELL-BASED CLOCK RECOVERY DEVICE

(57) Abstract

A clock recovery unit provides a clock recovery function in the receiving entity of a system to implement adaptation of constant bit-rate (CBR) services over an asynchronous transfer mode (ATM) or ATM-like network. Incoming cells are periodically sampled for buffer fill level L_i . The maximum fill level of undelayed cells L_{xj} is extracted from successive series of a predetermined whole number M of buffer-fill samples L_i . A frequency adjustment logic unit provides at its output a bit stream at a given clock frequency f_j . The frequency adjustment logic unit makes incremental adjustments to the clock frequency f_j tending to cause the steady state mean of the fill level L_{xj} , or its derivative, to move toward zero.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

- 1 -

CELL-BASED CLOCK RECOVERY DEVICE

This invention relates to digital transmission systems, and more particularly to a cell-based clock recovery (CBCR) device for providing a clock recovery
5 function in the receiving entity of a system to implement adaptation of constant bit-rate (CBR) services over an asynchronous transfer mode (ATM) or ATM-like network.

Asynchronous Transfer Mode (ATM)) is a high-speed digital communications protocol for which the basic
10 functional unit is a fixed-length 424-bit (53-byte) cell. Five bytes of each cell are allocated for routing and control, and the remaining 48 bytes are used for data transport. On entry to an ATM-based network, information is loaded into cells in accordance with standardized
15 formatting protocols called adaptation protocols. Once the cell has being filled, it is transported through the network as soon as possible. Buffers are used throughout ATM-based networks to deal with congestion, i.e., cases where more than one cell is ready for transport over a
20 given communications link.

As suggested by the name, constant bit rate (CBR) services transmit data bits at a nominally constant rate. More specifically, transmitting entities for CBR services use a reference clock to time the delivery of data bits.
25 Receiving entities must access the same clock, either directly or indirectly, to retrieve the data.

Reference clocks used in CBR services must operate within standardized specifications for mean operating frequency, maximum jitter and maximum wander. Jitter and
30 wander are constraints on high-frequency variability and low-frequency variability about the mean operating frequency, respectively. Jitter and wander are generally measured in terms of the offset of clock pulses from where they would be if the clock were consistently
35 operating at precisely its mean operating frequency.

- 2 -

Jitter is short-term deviation in the pulse center-points from where they would be if the clock was fixed at its mean frequency. Wander is the equivalent long term variation. Jitter is measured in the order of Hertz, whereas wander is measured in the order of hours or days.

The following specifications, drawn from the ANSI specification for T1 service, provide an example of requirements for CBR clocks. Here a UI is a unit interval or clock period, which for T1 is 648 ns:

- The mean operating frequency must be 1.544 ± 50 Hz.
- Jitter in the frequency band between 10 Hz and 40 KHz must be less than 0.5 UI peak-to-peak, and jitter in the frequency band between 8 KHz and 40 KHz must be less than 0.07 UI peak-to-peak.
- Wander must be less than 5 UI peak-to-peak over any 15-minute period, and must be less than 28 UI peak-to-peak over any 24-hour period.

Clock recovery is a process by which entities within a communications network gain access to a reference clock when needed. There are two basic approaches to clock recovery. The first approach entails applying relatively simple techniques to a clock signal that is separately transmitted to the receiving entity. The other approach is to extract the clock from an analysis of the periodicity of the received data signal. While the second approach is generally more complex to implement and more prone to error, it obviates the need for transmission of a separate clock signal. Cell-based clock recovery is an example of the second approach.

When CBR services are transported over an ATM network, a recommended basis for controlling the output clock in the receiver is to monitor the buffer fill level [CCITT, B-ISDN - ATM Adaptation Layer for Constant Bit Rate Services: Functionality and Specification, Draft

- 3 -

T1S1/92-570, November 13, 1992]. If the receiver's output clock is slower than the sender's input clock, then the buffer fill-level will tend to increase with time.

Conversely, if the receiver's output clock is faster than the sender's input clock, then the buffer fill-level will tend to decrease with time. The output clock can thus be adjusted based on trends in the buffer fill-level. In this arrangement, it is not necessary for both the sender and the receiver to have access to a common network clock. The approach can thus be used in more situations than the alternative recommended clock recovery method, i.e., the synchronous residual time stamp (SRTS) method.

Large coincident variations in the buffer fill-level can be expected. Firstly, the buffer fill-level plotted as a function of time looks like a saw-tooth because data are inserted into the buffer in cell-sized (53-byte) blocks but are drained from the buffer one bit at a time. Secondly, variation in the observed buffer fill level can be introduced by the sequencing and relative prioritization of tasks performed within the service adaptation system. For example, such processing jitter in the transmitting entity can cause variation in the time spacing between transmitted cells which will manifest itself as buffer fill-level variation in the receiving entity. Finally, time-varying queuing delays will occur at points of congestion within the network. As with processing jitter, such time-varying queuing delays manifest themselves as buffer fill-level variation.

Of the three sources of coincident fill-level variation mentioned above, the most problematic is queuing delays within the network. Processing jitter is under the control of the system designer and can be reduced to a manageable level by proper design. The saw-tooth effects can be minimized by roughly synchronizing

- 4 -

the sampling of the buffer fill-level with the arrival of ATM cells.

The Newbridge Mainstreet™ 36150 switch provides insight into the magnitude of fill-level variation due to queuing delays. The ATM cells pass through three switching stages, each of which has a 16-cell queue. The worst case scenario introduces a queuing delay of 48 cells. Given that 2.74 μ s (i.e. 53 bytes at 155 Mbits/sec) are needed to transmit an ATM cell, the queuing delay through a single switch can be up to 132 μ s. Allowing for processing jitter and the possibility of encountering a number of ATM switches, queuing delays can be expected to vary between zero and, say, 1000 μ s. For T1 service adaptation this translates to a buffer-fill-level variation in the order of +/-800 bits.

Other queues in the network can be much longer than those of the Mainstreet™ 36150 switch. For example, the Newbridge T3 line interface card has a queue for up to about 3 ms of data. A delay variation of up to 3 ms may be encountered if ATM cells for T1 service are passed through the T3 card.

An approximate analysis for satisfying the T1 jitter requirements is as follows: if the clock estimate is updated at a rate of about 20 Hz, the maximum allowable frequency mismatch is about 10 Hz to satisfy the 0.5 UI jitter upper limit. This corresponds to a clock period that is accurate to within about 4 ps. Larger errors in the clock period can in principle be tolerated if the frequency of the clock update is increased. However, this does not simplify the task because less new data is available from which to obtain each frequency estimate.

The magnitude of the task to be performed is quite striking. The requirement is to derive an unbiased estimate of the T1 transmitter's clock period to an accuracy of a few picoseconds by analyzing, in effect,

- 5 -

the periodicity of ATM cell arrival. However, cell arrival jitter of the order of $\pm 500 \mu\text{s}$ or more can be expected. It follows that the jitter variance must be reduced by a factor of in the order of 10^{12} .

5 Clock wander is important to control in CBR services and is a significant problem within conventional cell-based clock recovery systems. Wander is important because buffers within some CBR networks are sized based on the wander requirements, and may in certain circumstances
10 overflow if these requirements are not met. The problem with wander for conventional CBCR arises because the clock is, in effect, set based on low-pass filtered samples of the buffer fill level. Unfortunately, no matter how low one sets the roll-off frequency of the
15 low-pass filter, there is always in principle a lower-frequency component that can get through the filter. Such components appear as clock wander. Because the size of such lower-frequency components depends on a variety of uncontrolled factors, it is difficult to specify limits
20 of wander in CBCR systems and to verify conformance with standards for wander.

An object of the present invention is to address the aforementioned problems of the prior art.

25 According to the present invention there is provided a clock recovery unit for providing a clock recovery function in the receiving entity of a system to implement adaptation of constant bit-rate (CBR) services over an asynchronous transfer mode (ATM) or ATM-like network, comprising a buffer for receiving incoming cells, means
30 for periodically sampling the buffer fill level L_i , means for obtaining an estimate L_x of the buffer fill level on arrival of substantially undelayed cells from a series of buffer fill-level samples L_i , and a frequency adjustment logic unit providing at its output a control signal at a
35 given clock frequency f_j , said frequency adjustment logic

- 6 -

unit making incremental adjustments to said clock frequency f_j to cause the steady state mean of said estimate Lx_j of the buffer fill level on arrival of undelayed cells, or a derivative thereof, to move toward
5 a predefined optimal operating point L_{opt} .

In one embodiment, the estimate Lx_j is derived from blocks of fill-level samples, each block containing a predetermined number M of samples. However, the number of samples can change from block to block.

10 The control signal produced by the invention may be in the form of a bit stream, sine wave or other representation of the frequency of the derived clock.

The sampling of the buffer fill-level should normally be carried out in approximate synchronization
15 with the arrival of cells to minimize the effects of the saw-tooth shape of the fill level of the buffer for the reasons specified.

One estimator of the buffer fill level on arrival of undelayed cells is the maximum of a block of fill-level
20 samples. As indicated earlier the actual buffer fill level is in the form of a saw-tooth waveform because the cells arrive as a single block of 424 bits are then output one bit at a time at a constant rate. If the arriving cells are delayed, the fill level will tend to
25 fall because more bits will be output before new cells arrive. The maximum fill level will occur when the cells arrive on time. The estimated maximum fill level is thus representative of the fill level for undelayed cells.

While the interfering traffic in the network may
30 frequently create points of substantial congestion, it should also be relatively common for cells to pass through the network without substantial delay. Thus, the minimum of the cell delivery delays for a series of ATM cells should be relatively unaffected by the interfering
35 traffic. It follows that the maximum buffer fill-level

- 7 -

will also be relatively unaffected. Even if this is not completely true, it is reasonable to say that phenomena which increase the mean cell transmission delay will also increase the variance of the cell transmission delay, so
5 the minimum delay will undergo a smaller change than the mean delay or the maximum delay.

The buffer fill level on arrival of undelayed cells can also be estimated from the weighted sum of two or more statistics drawn from a block of fill-level samples.
10 For example, the mean fill level and/or the minimum fill level could be used in combination with the maximum fill level to obtain a composite estimate which, under certain conditions, leads to less wander than when the maximum fill-level is used in isolation.

15 When compared with the traditional approach of using the mean buffer fill level, the advantage of using the maximum buffer fill level or other estimate of the buffer fill level on arrival of undelayed cells, is particularly pronounced when a single bursty source of interfering ATM
20 traffic periodically swamps the capacity of some point of congestion within the network. In this case the mean buffer fill level observed by the receiving entity will be severely affected by the interfering traffic, but the maximum buffer fill level or other estimate of the buffer
25 fill level on arrival of undelayed cells will be relatively unaffected.

The invention also provides a method of providing a clock recovery function in the receiving entity of a system to implement adaptation of constant bit-rate (CBR)
30 services over an asynchronous transfer mode (ATM) or ATM-like network, characterized in that it comprises the steps of receiving incoming cells in a buffer; periodically sampling the buffer fill level L_i ; estimating L_x of the buffer fill level on arrival of substantially
35 undelayed cells from a series of buffer fill-level

samples L_i ; outputting a control signal at a given clock frequency f_j ; and making incremental adjustments to said clock frequency f_j to cause the steady state mean of the estimate Lx_j of the buffer fill level on arrival of
5 undelayed cells, or a derivative thereof, to move toward a predetermined optimal operating value.

The invention has been described with reference to an ATM network, but it is applicable to any similar type of packet-switched network having cells of data that are
10 propagated through the network. The cells need not necessarily be of fixed length, and the invention is equally applicable to a packet-switched network employing blocks of data of variable size.

The invention will now be described in more detail,
15 by way of example only, with reference to the accompanying drawings, in which:-

Figure 1 is a general block diagram of a cell-based clock recovery unit in accordance with the invention;

Figure 2 is a general block diagram of one
20 embodiment of a clock recovery unit in accordance with the invention;

Figure 3 illustrates a mathematical model of the clock recovery method employed in the invention; and

Figure 4 shows the variation in buffer bill level
25 with time.

Referring now to Figure 1, incoming 53-byte ATM cells 21 are input to buffer 1 as they arrive from the network. The buffer outputs bits at a constant bit rate. As a result the fill level of the buffer 1 as a function
30 of time can be represented by a saw-tooth waveform as shown in Figure 4, which shows the fill level for undelayed cells in solid outline. The rising edge occurs when a cell arrives. If the cells are delayed the maximum

- 9 -

fill level falls as shown in broken outline because more bits are output before the arrival of the next cell.

Referring again to Figure 1, the fill-level L_i of the buffer 1 is monitored by buffer fill-level detector 2.

5 The samples of the buffer fill-level L_i are obtained by periodic sampling of the buffer fill-level in approximate synchronization with the arrival of ATM cells to minimize the effects of the sawtooth shape of the fill level waveform.

10 The samples L_i are passed to block 3, which is an estimator of buffer fill level for undelayed cells. This extracts its estimates from a sequential number M of buffer fill-level samples L_i , producing Lx_j , which is only updated after M new samples of L_i have been scanned. That
15 is subscript j increments M times slower than subscript i . M is a predetermined configuration parameter for the invention.

The fill level estimator can be a local maximum estimator, for example, or alternatively a unit taking
20 the weighted summation of the mean fill level, the maximum fill level, or other statistics from a block of fill-level samples.

The output of block 3 is then fed to subtractor 4, which subtracts a pre-determined steady state optimum
25 buffer fill level L_{opt} from Lx_j , and the result is passed to the frequency adjustment logic block 6 along with an estimate of the derivative of Lx_j produced by differentiator 5.

The choice of L_{opt} is a choice between starvation
30 avoidance and delay minimization. Larger values of L_{opt} are needed to avoid starvation (buffer underflow) when cells are severely delayed during transmission. Small values of L_{opt} are preferred to keep the mean delay within the network low. The parameter L_{opt} provides a direct
35 means of obtaining a balance between these two factors.

- 10 -

The frequency adjustment logic block 6 generates a signal representing a clock frequency f_j . Frequency adjustment logic block 6 then makes small incremental adjustments to the clock frequency f_j such that the steady-state mean of one or both of its inputs tends toward zero.

Figure 2 provides a detailed illustration of a preferred embodiment of the invention. L_i represents sample i of the buffer fill level, L_{\max_j} sample j of the recovered clock frequency, L_{opt} the optimum steady-state maximum buffer fill level, is the index for the samples of buffer fill level, and j is the index for samples of the clock frequency and the maximum buffer fill level.

The Divide-by-M block 9, Maximum Extractor Block 10 and the Maximum Sample-and-Reset block 11 correspond to the block 3 of Figure 1 since they provide an estimate of the fill level of undelayed samples. The block 10 outputs a signal representing the maximum of the fill level samples L_i received from the fill level detector 2. Block 11 outputs the maximum L_{\max_j} for each M samples and at the same time resets the maximum extractor 10. The result is an output signal L_{\max_j} that represents the buffer fill level for undelayed samples.

An optimization signal L_{opt} , which is an optimization parameter for the system that depends on the operating conditions, is subtracted from L_{\max_j} in summer 4. The output of summer 4 is input to the J sample delay block 12 and summer 13, which together correspond to the differentiator block 5 of Figure 1.

Multipliers 14, 15 and summer 16 correspond to the Frequency Adjustment Logic block 6 of Figure 1.

The embodiment shown in Figure 2 implements the recursion relationship:

$$f_i = f_{j-1} + \alpha \times HDLIM_{\alpha} (L_{\max_j} - L_{\max_{j-j}}) + \beta \times (HDLIM_{\beta} (L_{\max_j} - L_{\text{opt}}))$$

- 11 -

where

f_j = output clock frequency (Hz)

L_{max_j} = the maximum of M successive samples of the buffer fill level (bits)

5 M = the length of the block from which each L_{max_j} is extracted (samples)

L_{opt} = the optimal buffer fill level (bits)

α = first-order adaptation factor (Hz/bit)

β = second-order adaptation factor (Hz/bit)

10 J = the block separation

$HDLim_\alpha$ = hard limit with threshold THR_α

$HDLim_\beta$ = hard limit with threshold THR_β

In the absence of coincident sources of variation, the buffer fill-level changes in direct proportion to the
15 difference between the transmitting entity's clock and the receiving entity's clock. Thus, α serves to adapt f_j toward the input clock frequency. The other factor β performs the secondary function of moving the buffer fill-level toward the preselected "optimal" value L_{opt} .
20 Without β , the unit will come to rest at an arbitrary fill level which may be too low to avoid starvation (buffer underflow) or too high to avoid buffer overflow. Furthermore, without β the end-to-end transmission delay of the adaptation system is not controlled.

25 The hard limiting operator $HDLIM_\alpha$, serves to reduce the impact of occasional large spikes in the driving function that can arise when the network delay suddenly changes. The invention relies on what amounts to the derivative of the buffer fill level. A step in the fill
30 level thus translates to a large spike in the clock adjustment. Such steps in the fill level can result if the network delay properties suddenly change.

The hard limiting operator $HDLim_\beta$ serves to control the excursion the output frequency when a large change in
35 the buffer fill level is needed. Such large changes can

- 12 -

occur if the network delay suddenly changes or if the target fill level (L_{opt}) is altered.

The following is a general analytical analysis of the embodiment shown in Figure 2, which will explain the underlying principle of the invention in more detail.

A Z-Transform representation of a generalized version of the recursion relationship presented above is:

$$F(z) = D(z) (L(z) + N(z))$$

where

10 $F(z)$ = Z transform of the output frequency
 $D(z)$ = Z transform of the adaptation filter
 $L(z)$ = Z transform of the maximum buffer fill level
 $N(z)$ = Z transform of the error in estimates of $L(z)$

15 The maximum buffer fill level satisfies the following relation in the absence of jitter provided that the frequency parameters remain approximately constant over the interval of analysis:

$$L_{max_j} = L_{max_{j-1}} + T (f_{in_{j-1}} - f_{j-1})$$

20 where

$f_{in_{j-1}}$ = the frequency of the CBR source clock
 T = the time separation of fill-level estimates L_{max_j} and L_{j-1}

25 By deriving the Z transform of L_{max_j} and substituting it into the initial expression for $F(z)$ we obtain:

$$F(z) = H(z) F_{in}(z) + ((H(z)(z-1))/T) N(z)$$

where

$$H(z) = D(z)T/(z - 1 + D(z)T)$$

30 The above expressions are represented by a digital phase-locked loop (DPLL)-like structure as shown in Figure 3. Since the element $1/(z-1)$ is an integrator, the buffer fill-level is the integral of the frequency

- 13 -

difference, and thus is effectively the phase error of the loop. The linearized baseband model of a classical DPLL can be derived, for example as described in "A survey of Digital Phase-Locked Loops," Proceedings of the IEEE, April 1981, pp 410-431, by backing the integrator $1/(z-1)$ out of the loop through the summer. An equivalent circuit for Figure 3 is obtained by removing the $1/(z-1)$ block from its current position in the loop and applying it instead to both f_j and f_{in} , prior to their summation.

The main remaining difference between this result and that described in the above article is in scaling of the noise and the loop filter. Thus, much of the standard analyses commonly performed for classical DPLLs are applicable.

The analytical representation presented above has a loop filter of the form:

$$D(z) = (\alpha(1-z^{-J}) + \beta)/(1-z^{-1})$$

from which is obtained:

$$H(z) = \frac{\alpha T z^{-1} + \beta T z^{-1} - \alpha T z^{-1-J}}{1 - (2 - \alpha T) z^{-1} + z^{-2} + \beta T z^{-1} - \alpha T z^{-1-J}}$$

For T1 service adaptation the time interval T is:

$$T \approx BM/f_{in}$$

where

B = number of data bits per cell (376 for ATM)

M = number of cells per block

f_{in} = nominal source clock frequency (1.544 Mhz for T1)

The stability of the above-described system will now be considered. Through repeated application of the Jury stability test to $H(z)$ when the loop filter has the form

$D(z) = (\alpha(1-z^{-J}) + \beta)/1-z^{-1}$, it was found that the loop is stable if the following conditions are satisfied:

- 14 -

$$0 < \alpha T < 2^{2-J} - \beta T/2$$

$$0 < \beta T < 2^{3-J} - 2\alpha T$$

The loop will exhibit a highly oscillatory behavior as it nears the limits of stability, and extra delay in the loop can push it towards instability. It is thus
 5 important to operate well inside the stability limits.

Extra delays in the loop lead to an increased tendency for oscillatory behavior, and can make an otherwise stable loop unstable. For example, if $J = 1$ and
 10 if there is an extra one-sample delay in the loop (i.e., a delay of T seconds), the Jury stability test produces the following limits for obtaining stability:

$$0 < \alpha T < .5 + \sqrt{.25 + \beta T}$$

$$0 \leq \beta T < \alpha T(1 - \alpha T) \quad \text{when } J = 1 \text{ and delay} = T$$

One can see that the upper limits are substantially
 15 less than they would be without the extra delay.

A properly-designed loop must provide more than just stability: it should have few if any oscillations when presented with sudden input changes. For the present application, such oscillations occur when β is too large
 20 relative to α . In the special case of $J = 1$, the transfer function becomes:

$$H(z)|_{J=1} = \frac{2(1-a)z - (1-a^2-b^2)}{(z-a)^2 + b^2}$$

where

$$a = 1 - \frac{(\alpha T + \beta T)}{2}$$

$$b^2 = \beta T - \frac{(\alpha T + \beta T)^2}{4} = \beta T - (1-a)^2$$

From an inspection of the denominator of $H(z)|_{J=1}$ it is apparent that the loop response is critically damped
 25 when $b = 0$. This means that the following inequality should be applied to avoid oscillatory behavior:

$$\beta T \leq 1 - \alpha T - 2\sqrt{1 - \alpha T} \quad \text{when } J=1$$

- 15 -

Now, with $\beta = 0$ the convergence properties of the loop are quite consistent when $\alpha \times J$ is constant. The choice of J also has little effect on the influence of β on the loop dynamics. Therefore, the following expression
 5 can be used to obtain a nearly-critically-damped loop for arbitrary J :

$$\beta T \leq 2 - \alpha J T - 2\sqrt{1 - \alpha J T}$$

It should be emphasized that this expression is only an approximation. For example, when $J > 1$ it is possible
 10 for α alone to introduce oscillations. Additional loop delays can also introduce oscillations. Thus, it is prudent to choose βT such that the loop is somewhat overdamped.

A simple measure of the convergence rate of the loop
 15 can be derived from the proportion of the estimated frequency difference that is corrected in a single step. While this ignores interaction between α and β and provides only a tangential approximation for the convergence, it has been found to be consistently
 20 representative for all practical parameter selections. Now, we know that

$$L_{\max_j} - L_{\max_{j-1}} = TJ(f_{in_j} - f_j)$$

From this it appears that the α that would correct a frequency mismatch in a single step is:

$$25 \quad \alpha_1 = 1/TJ$$

Thus, the proportion of frequency correction in a single step is α/α_1 , and rate of frequency change per unit time is:

$$\text{RATE} = \alpha/\alpha_1 T = \alpha J$$

30 The interpretation of RATE is that a step discontinuity in the input frequency will be resolved by the loop in about $1/\text{RATE}$ seconds. The actual rate tends to be slower. For example, if one inspects the actual

response of a critically-damped loop to a frequency step, one finds that the time it takes for the output frequency to match the input (i.e., the first zero crossing in the frequency error plot after the step discontinuity), is
 5 about 1.8/RATE. It is interesting to note that the convergence rate is not affected by the block size (M) or, implicitly, the update period T.

There is a direct nonlinear relationship between RATE and the -3 dB rolloff point of the closed-loop
 10 frequency response of the loop (f_{3dB}). If RATE is well below its maximum, i.e., if $\alpha \ll \alpha_1$, then $f_{3dB} \approx \text{RATE} / 6$ Hz. The divisor is larger for higher values of RATE. The divisor is between 5.5 and 6 for the configurations recommended in this document.

15 An approximate analysis of worst-case buffer fill-level excursion during frequency convergence is to temporarily set β to zero. This is reasonably accurate because the approximate behavior of a critically damped second-order loop when presented with a step
 20 discontinuity in the input frequency is to first resolve the frequency mismatch at the rate determined by α and J, and then slowly bring back the buffer fill-level to L_{opt} at a rate determined by β . Now, since the proportion of frequency correction on each iteration is α/α_1 , the
 25 change in fill level over N iterations is:

$$\Delta L = \Delta f T (1 + (1 - \alpha/\alpha_1) + (1 - \alpha/\alpha_1)^2 + (1 - \alpha/\alpha_1)^{N-1})$$

where

ΔL = fill-level change (bits)

Δf = size of the frequency step (Hz)

30 Taking the limit as N goes to infinity one obtains:

$$\Delta L = (\Delta f T \alpha_1) / \alpha = \Delta f / \text{RATE}$$

Simulations indicate that this expression overestimates the fill-level excursion by a factor of

- 17 -

about 1.3 for a critically-damped loop. The overestimate is less for over-damped loops.

An analysis will be made of the frequency disturbance due to fill level adjustment. This is an analysis of the peak frequency disturbance when the loop acts to resolve a deviation of the buffer fill-level from its prespecified "optimal" value (L_{opt}). Such a response will result if the L_{opt} parameter is adjusted during processing or if clock recovery device is activated when the buffer fill-level deviates significantly from L_{opt} . For the present approximate analysis, $L_{max_i} - L_{opt}$ in the adaptation equation are replaced with the constant ΔL_{opt} , the hard-limits are disabled, and the resulting steady-state frequency difference Δf is computed. If ΔL_{opt} is the initial size of the fill-level deviation, then Δf will be a rough estimate of the peak frequency disturbance. The modified adaptation equation comes to rest when $\alpha(L_i - L_{i-J}) = \beta L_{opt}$. Since $L_{max_i} - L_{max_{i-J}} = TJ(f_{in} - f_{out})$, it follows that the steady-state frequency difference is:

$$\Delta f_{ss} \approx \frac{\beta \Delta L_{opt}}{\alpha TJ} = \frac{\beta \Delta L_{opt}}{RATE \times T}$$

This expression overestimates the peak frequency disturbance because it ignores the fill-level adaptation that occurs before we reach the peak. Simulations indicate that it overestimates by a factor of about 1.3 for a critically-damped loop. The overestimate is less for over-damped loops.

It is also necessary to consider the hard limiting effects of THR_p and THR_a . The hard-limiting threshold THR_p is present to control the range of frequency excursion when the fill-level deviates from its prespecified "optimal" value. Using the same logic as was presented in the preceding section, it can be shown that a sustained fill-level deviation which exceeds THR_a will result in the following frequency offset:

$$\Delta f_{ss} \approx \frac{\beta THR_{\beta}}{\alpha TJ} = \frac{\beta THR_{\beta}}{RATE \times T}$$

The hard-limiting threshold THR_{α} limits the influence of sudden large changes in the fill level due possibly to extraordinary amounts of error in the fill level

5 estimates. This is intended for use as a means of rejecting widely deviant samples. One should avoid bringing it too close to the normal range of variation because it reduces α and may thus push the loop into oscillation or instability when β is nonzero.

10 Formal analysis of the noise output of the loop requires a rather tedious contour integration and is not needed for gathering a general insight into its behaviour. First, let $\beta = 0$ for the purpose of this analysis, because it is always much less than α for
15 practical loops. This leads to the following approximation for the loop filter:

$$f_j - f_{j-1} \approx \alpha(L_{\max_j} - L_{\max_{j-1}})$$

If we run through the recursion we find that each output sample (f_j) is the sum of J successive input
20 samples L_{\max_j} . If it is assumed that the input noise is uncorrelated, then

$$\sigma_f^2 = \alpha^2 J \sigma_L^2$$

where

σ_L^2 = variance of the maximum buffer fill level
25 (bits²)

σ_f^2 = variance of the output frequency (Hz²)

For the purpose of this analysis it is assumed that σ_L^2 is the same as the variance of the average buffer fill level obtained from a block of M fill-level samples.

30 Thus,

$$\sigma_L^2 = \frac{fin^2 \sigma_f^2}{M}$$

- 19 -

where

σ_L^2 = variance of the ATM cell arrival times (sec²)

This leads to the following expression for the frequency jitter (in Hz²)

$$\sigma_f^2 \approx \frac{\alpha^2 J \text{fin}^2 \sigma_\delta^2}{M} = \frac{\text{RATE} \alpha \text{fin}^2 \sigma_\delta^2}{M}$$

Finally it is common to specify jitter in terms of unit intervals rather than frequency, where a unit interval is a period of the T1 clock (648 ns).

The following formula, with the dimension UI² provides a rough comparison with such specifications:

$$\sigma_{UI}^2 = \sigma_f^2 T^2 \approx \text{RATE} \times B^2 M \alpha \sigma_\delta^2$$

The choices of M and J are affected by factors other than noise reduction. For example, each increment of J increases the order of the loop filter, thereby making it more difficult to analyze and tightening the stability constraints. A small M implies a small T which leads to a potentially larger impact of the delay on the loop performance and stability.

The preferred settings for T1 clock adaptation are summarized in the Table below. For both modes the loop is updated with a frequency of about 20Hz, i.e. $T=200/4106 = 0.0487$ secs. The loop parameters are well within the stability range bounds. The fast adaptation configuration produces a system having a -3dB rolloff of about 0.1 Hz. It produces a maximum fill level excursion of $\Delta L < 200$ bits in response to a 100Hz frequency step, and a maximum disturbance of $\Delta f_{ss} < 30\text{Hz}$ in response to a large fill-level adjustment. The slow adaptation configuration produces a system having a -3dB rolloff of about 0.03 Hz. It produces a maximum fill level excursion of $\Delta L < 500$ bits in response to a 100Hz frequency step, and a maximum disturbance of $\Delta f_{ss} < 11\text{Hz}$ in response to a large fill-level adjustment. It is finally both configurable and

- 20 -

sufficiently overdamped to provide a maximum low frequency amplification of less than 0.5 dB.

Table

Parameter	Adaptation	Mode
	Fast	Slow
RATE	1/2	1/5
α	0.0625	0.025
β	$\alpha/70$	$\alpha/200$
J	8	8
M	200	200
THR _{α}	256	256
THR _{β}	800	800

- 5 The described embodiments of the invention are capable of providing CBR service adaptation in an ATM network.

- 21 -

Claims:

1. A clock recovery unit for providing a clock recovery function in the receiving entity of a system to implement adaptation of constant bit-rate (CBR) services over an asynchronous transfer mode (ATM) or ATM-like network, characterized in that it comprises a buffer for receiving incoming cells; a sampling unit for periodically sampling the buffer fill level L_i ; an estimation unit for obtaining an estimate Lx_j of the buffer fill level on arrival of substantially undelayed cells from a series of buffer fill-level samples L_i ; and a frequency adjustment logic unit providing at its output a control signal at a given clock frequency f_j , said frequency adjustment logic unit making incremental adjustments to said clock frequency f_j to cause the steady state mean of said estimate Lx_j of the buffer fill level on arrival of undelayed cells, or a derivative thereof, to move toward a predefined optimal operating point.
2. A clock recovery unit as claimed in claim 1, characterized in that said estimation unit comprises means for extracting the maximum fill level of said predetermined number of samples.
3. A clock recovery unit as claimed in claim 1, characterized in that said estimation unit comprises means for deriving a composite estimate from the mean fill level in combination with the maximum fill level.
4. A clock recovery unit as claimed in claim 1, characterized in that said estimation unit comprises means for deriving a composite estimate from the minimum fill level in combination with the maximum fill level.
5. A clock recovery unit as claimed in claim 1, characterized in that said sampling unit samples the buffer fill-level in approximate synchronization with the arrival of ATM cells.

- 22 -

6. A clock recovery unit as claimed in claim 1, further comprising a subtractor for subtracting a predetermined optimal value L_{opt} from the estimate of the buffer fill-level Lx_j to produce a difference value, a differentiator for
5 deriving the derivative of said estimate of the buffer fill-level Lx_j , said difference value and the derivative of said estimate of the buffer fill-level Lx_j being applied as first and second inputs to said frequency adjustment logic unit which is adapted to incrementally adjust said clock
10 frequency f_j to cause one or both of the inputs thereof to move toward a said predetermined optimal value.

7. A clock recovery unit as claimed in claim 2, characterized in that said estimation unit extracts the maximum fill level L_{max_j} and includes a divider for dividing
15 the clock rate of the samples by M and a Maximum-Sample-and-Reset unit for producing at its output a signal representing L_{max_j} .

8. A clock recovery unit as claimed in claim 7, characterized in that said differentiator is provided by a J
20 sample delay circuit and a subtractor.

9. A clock recovery unit as claimed in claim 8, characterized in that said frequency adjustment logic unit includes a pair of multipliers each having respective first inputs receiving a value dependent on the maximum fill-level
25 L_{max_j} and the derivative thereof, and second inputs receiving predefined configuration parameters α and β .

10. A clock recovery unit as claimed in claim 9, characterized in that said frequency adjustment logic unit comprises hard limit units connected to respective
30 multipliers, and a summer for summing the outputs of said multipliers to produce said control signal, said summer having its output looped back to a third summation input.

11. A clock recovery unit as claimed in claim 10, further comprising a one sample delay unit between the output of
35 said summer and said third input thereof..

- 23 -

12. A clock recovery unit as claimed in claim 11, characterized in that α is about 0.0625 and β is about 0.0009 in a fast adaptation mode.

13. A clock recovery unit as claimed in claim 12,
5 characterized in that α is about 0.025 and β is about 0.0001 in a slow adaptation mode.

14. A clock recovery unit as claimed in claim 1,
characterized in that said means for obtaining an estimate
of the buffer fill level obtains said estimate from
10 successive series of samples, each having a number M of samples.

15. A clock recovery unit as claimed in claim 14,
characterized in that M is constant from one series of
samples to the next.

16. A clock recovery unit as claimed in claim 14,
15 characterized in that M varies from one series of samples to the next.

17. A method of providing a clock recovery function in the
receiving entity of a system to implement adaptation of
20 constant bit-rate (CBR) services over an asynchronous transfer mode (ATM) or ATM-like network, characterized in that it comprises the steps of receiving incoming cells in a buffer; periodically sampling the buffer fill level L_i ; estimating the buffer fill level on arrival of substantially
25 undelayed cells Lx_j from a series of buffer fill-level samples L_i ; outputting a control signal at a given clock frequency f_j ; and making incremental adjustments to said clock frequency f_j to cause the steady state mean of the estimate Lx_j of the buffer fill level on arrival of
30 undelayed cells, or a derivative thereof, to move toward a predetermined optimal value.

18. A method as claimed in claim 17, characterized in that said estimate Lx_j is derived from the maximum buffer fill level.

- 24 -

19. A method as claimed in claim 17, characterized in that M is constant from one series of samples to the next.

20. A method as claimed in claim 17, characterized in that M varies from one series of samples to the next.

1/3





Figure 1

2/3



3/3




Figure 3




Figure 4

INTERNATIONAL SEARCH REPORT

In. ional Application No
PCT/CA 95/00320

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04J3/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO,A,88 07297 (BELL TELEPHONE MFG ;ALCATEL NV (NL)) 22 September 1988 see page 3, line 26 - line 31 see page 4, line 32 - page 5, line 9 see page 6, line 34 - page 7, line 5 see page 7, line 35 - page 14, line 27 ---	1,17
A	SINGH R P ET AL 'Jitter and clock recovery for periodic traffic in broadband packet networks', IEEE TRANSACTIONS ON COMMUNICATIONS, MAY 1994, USA, VOL. 42, NR. 5, PAGE(S) 2189 - 2196 , ISSN 0090-6778 see page 2191, column 2 - page 2193, column 2, paragraph III --- -/--	1,17

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

31 August 1995

Date of mailing of the international search report

1 8. 09. 95

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+ 31-70) 340-3016

Authorized officer

Van den Berg, J.G.J.

INTERNATIONAL SEARCH REPORT

Int ional Application No
PCT/CA 95/00320

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>I.B.M. TECHNICAL DISCLOSURE BULLETIN, vol. 37, no. 48, April 1994 ARMONK (US), pages 459-463, XP 000451314 NOMEN NESCIO 'Rate based end-to-end clock synchronisation.' see figure 2</p> <p style="text-align: center;">-----</p>	1,17

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/CA 95/00320

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO-A-8807297	22-09-88	BE-A- 1000415	22-11-88
		AU-B- 607475	07-03-91
		AU-A- 1395288	10-10-88
		DE-A- 3875484	26-11-92
		EP-A, B 0365526	02-05-90
		JP-T- 2502776	30-08-90
		US-A- 5027351	25-06-91
